A. Overview Information

I. Federal Agency Name: Defense Advanced Research Projects Agency,
II. Information Processing Technology Office (DARPA/IPTO)
III. Title: Bootstrapped Learning (BL)
IV. Announcement Type: Request for Information (RFI)
V. Solicitation Number: RFI 06-33
VI. CFDA Number: 12.910
VII. Key Dates
   a. DARPA posting of specification draft: September 11, 2006
   b. RFI Registration and Workshop Statement of Interest Deadline: 12:00 NOON (ET), September 15, 2006
   c. RFI Response Due: 12:00 NOON (ET), September 20, 2006
   d. Workshop Registration Deadline: September 20, 2006
   e. Workshop: September 26, 2006 – Palo Alto, CA
   f. Workshop Sidebar/Teaming Sessions -- Optional (one-on-one with PM): September 27, 2006 – Palo Alto, CA

B. Full Text of Announcement

I. Description

In accordance with FAR 35.007(j), the Information Processing Technology Office (IPTO), Defense Advanced Research Projects Agency (DARPA) requests information on the technical approach of an anticipated research program on “Bootstrapped Learning.” Specifically, the Request for Information (RFI) is interested in natural instruction methods, interaction language specification, interlingua specification, suggestions for curricula and intellectual property considerations. These five areas are discussed in detail in the “Request for Information (RFI) Section” below. DARPA currently anticipates that the program will occur in three phases over a period of three years. DARPA expects to fund ‘knowledge engineers’ who will develop very rich testing domains, and several ‘learning teams’ devoted to basic Artificial Intelligence (AI) research. The “learning” teams will be composed of significant academic and industrial components. A Broad Agency Announcement (BAA) and/or other solicitation, regarding this program, is anticipated later this year.

Background

Bootstrapped Learning (BL), as envisioned by the anticipated program, tackles the problem of creating an electronic student – a computational system that can learn by being taught through the use of methods related to human-to-human instruction. This represents a break from traditional machine learning (ML) in that, in BL, the assumption is that there is a (human) instructor who, a priori, has the knowledge and capabilities that the student is expected to learn. Traditional ML, by contrast, primarily focuses on learning capabilities that are not possessed by the instructor (user of the technology).

In some ways, ML can be seen as the problem of knowledge discovery, and BL as the problem of knowledge *transfer*. The ability for humans to transfer knowledge to (i.e., instruct) a machine in similar ways to how they transfer knowledge to other humans would have a profound effect on many AI systems---it would allow ordinary people (i.e., non-AI experts) to enter knowledge into systems and change system behavior. Table 1 compares bootstrapped learning and traditional machine learning along several dimensions.

Bootstrapped Learning	**Traditional Machine Learning**
Knowledge transfer	Knowledge discovery
Learn by being taught	Learn from data
Requires small sets of data (examples)	Requires large sets of data
Instructor knows the complete answer	No one knows the complete answer
"Laddered" curricula (lessons are structured to teach base knowledge first; later lessons build upon earlier lessons)	Unstructured learning

Table 1: Comparison of Bootstrapped Learning and Traditional Machine Learning

Bootstrapped learning is not limited to the types of instruction that take place in a classroom setting. Rather, it comprises the disparate ways that knowledge is transferred between humans. There are many ways that human instructors interact with students. This will be termed *natural instruction (NI) methods*. An example of an NI method is *By Annotated Example*, in which an instructor provides a (small) set of examples (e.g., photographs) that illustrate the concepts to be learned, possibly commenting on them using natural language and gesture to point at salient features of the examples. Another example is *By Demonstration* in which the instructor demonstrates the procedure to be learned, such as when a swimming instructor demonstrates a stroke. The final example is *By Practice*, where the student is given a context and goal to practice against in order to refine skills. These NI methods are described in more detail below.

The goal of bootstrapped learning is to create *domain-independent* learning algorithms for NI methods. As mentioned above, these are very different research problems than traditional machine learning has tackled. At first glance, learning from the *By Annotated Example* method may seem very much like a supervised learning task. There are some important differences: the NI method provides only as much training data as would typically be shown to a human student, usually on the order of tens of examples, and not millions (or even thousands) as in the case with ML. However, since *By Annotated Example* is intentional instruction, one can assume that the small set of examples was not randomly sampled, as is the assumption for traditional ML. Rather, it can be assumed that the given examples were carefully chosen by the instructor in order to guide learning. Because instructors are assumed to understand the concept being learned, they can provide many forms of hints that simplify the learning task. In general, any information that one human might provide to another during instruction could be part of the input to a given NI method.

Bootstrapped learning also reduces learning complexity (over traditional machine learning) by the use of “laddered” curricula, where an instructor has carved up the learning task into lessons or “rungs,” which represent smaller, much more constrained, learning tasks. A student can then “climb” the “learning ladder” by learning one rung at a time.

Anticipated Program Structure

The goal of the anticipated Bootstrapped Learning program is to foster fundamental research in domain-independent BL. As such, the anticipated structure of the program (as shown in Figure 1) is as follows: several teams (*learning system builders*) will identify a useful subset of natural instruction (NI) methods, and build computational models of a student that can learn from that fixed subset of NI methods. In order to ensure the generality of those systems, another team of performers (*curriculum developers*) will create learning curricula in disparate domains on which the learning system builders will be tested. It is important to emphasize that the learning teams will not know what these domains are beforehand. In addition, in at least one domain, the curriculum team will compare the performance of the learning systems to human learning on the same curriculum.

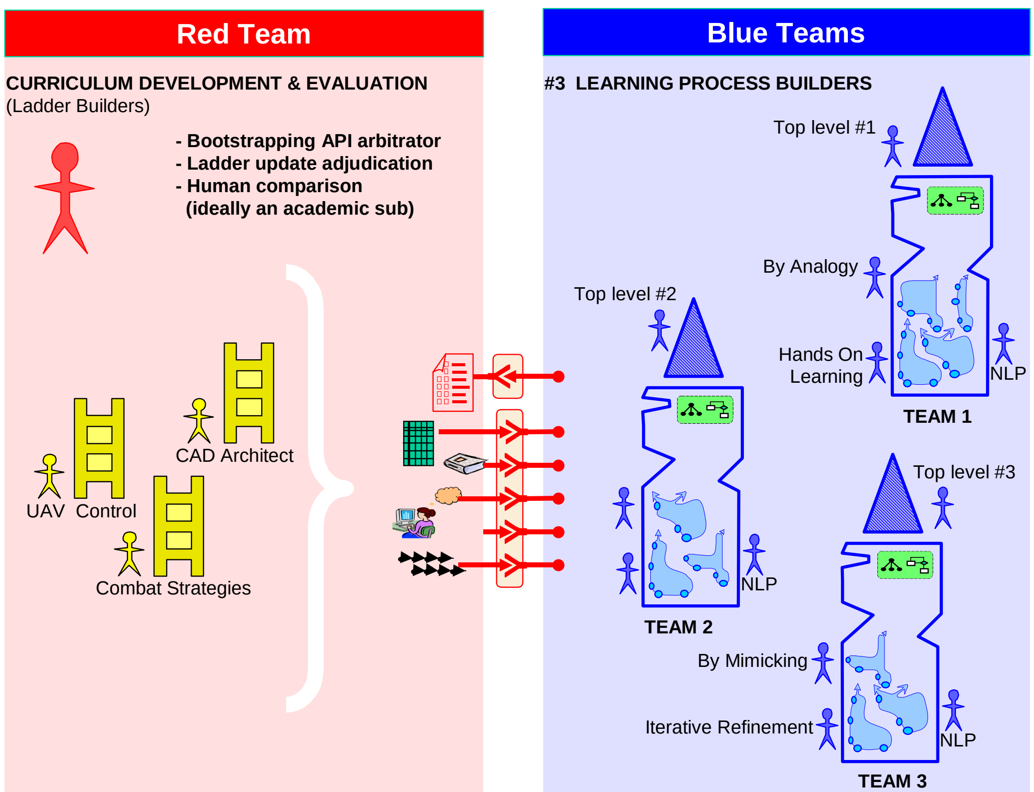

Figure 1: Bootstrapped Learning Program Structure

In order for this to work, it will be necessary to have a common interaction language (represented in Figure 1 as the lines between the two types of teams) in which all lessons in a curriculum will be encoded by the curriculum developers, and understood by the learning teams. This language will include specifications for both *interaction modalities* (e.g., {constrained} natural language, gesture, perception, and diagrams) as well as natural instruction methods that use those modalities. For example, a potential protocol for the NI method of *By Annotated Example* might be a list of world snapshots (perception) with optionally-attached linguistic explanation (relevant sub-expressions), and gesture (pointing) at salient features of the snapshots. The NI method *By Practice* might include a set of practice problems (world setups) and a specified performance goal expressed either linguistically or during previous learning, which would then allow the student to practice the given task, to maximize performance on the goal.

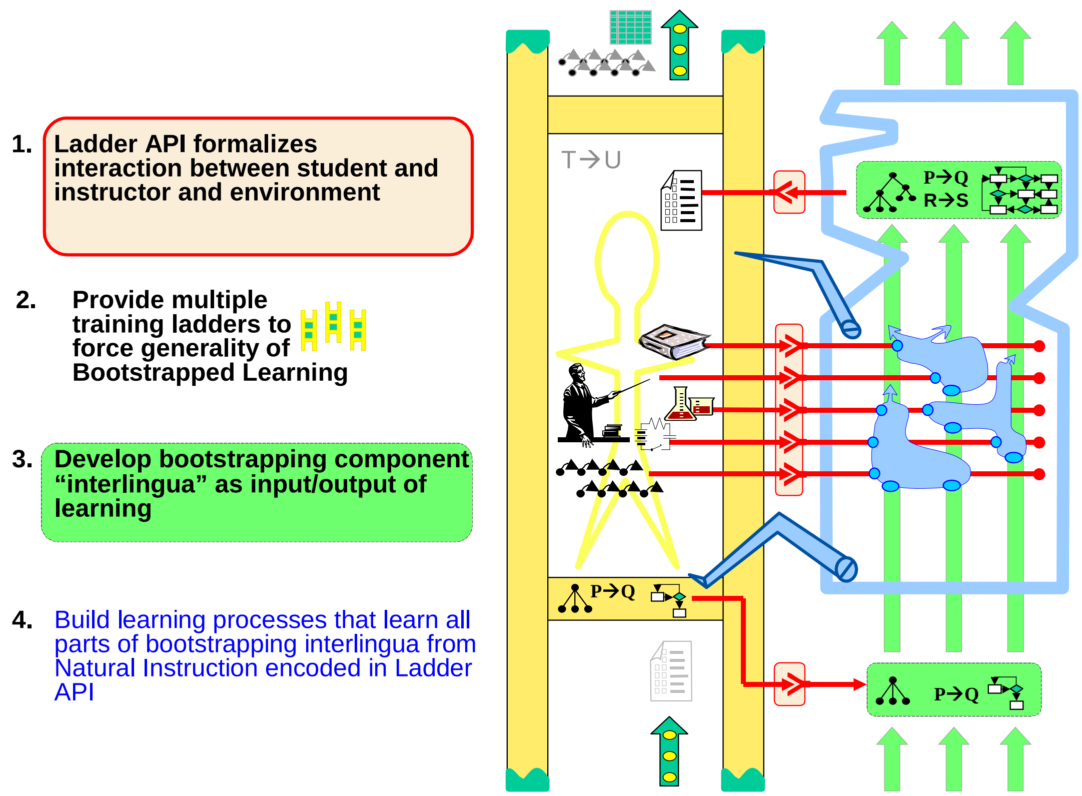

Figure 2: The General Program Approach

Figure 2 shows more detail of the general approach of the program by focusing in on a single curriculum, or ladder. In addition to a common interaction language between ladders and learners, it will be necessary to have a common "interchange" representation of learned knowledge, which will be termed the *interlingua*. As mentioned above, students will first be taught simple concepts and later lessons will depend on knowledge learned in earlier lessons. This means that learning components must generate knowledge learned in a form in which other learning processes can access it for future learning. There will also be certain NI methods (such as *By Practice*) that will have the purpose of *refining* previously learned knowledge. In order to do so, the previously learned knowledge must be represented in a way that allows easy modification by learning processes.

Additionally, the student's assumed prior knowledge must also be encoded in the interlingua. Two types of prior knowledge are distinguished: *genetic knowledge* and *injected knowledge*. Genetic knowledge is the knowledge assumed by *any* learner in *any* domain, and is public knowledge in the sense that it is known to both the (human) curriculum developers and the (human) learning system developers. In some sense, genetic knowledge can be seen as the general contract between curriculum and learning system builders. Genetic knowledge will include, among other things, the specification of the interaction language, the NI methods, and the interlingua, as well as a general upper-level ontology into which domain-specific knowledge will be placed.

Injected knowledge (shown as the bottom rung of the ladder in Figure 2) reflects the need for a curriculum to make assumptions of a student's prior knowledge *within a particular domain*. Although, in general, it should be possible for bootstrapped learning to proceed from only genetic knowledge, in practice, one may want to start at a higher level in order to concentrate on more interesting problems in the domain. To support this, the first rung of each ladder may "inject" the prior domain knowledge (encoded in interlingua) that is required to attain the next rung into each learner. The status of this knowledge will then be the same as if the learner had actually learned it from a curriculum.

The following paragraphs briefly describe the different expected performer roles, a description of the interaction language, NI methods, and the interlingua in more detail. Also discussed is an example of a lesson and the expected learning to be achieved from it

Learning System Builder Teams

The program anticipates funding several teams that will build learning systems. Each team will construct a self-contained learning system that is capable of learning from a curriculum that uses a given set of NI methods. Each team will have an *integrator*, who will create the student infrastructure and top-level control strategy. Integrators will require experience integrating large AI programs as well as a depth of understanding in machine learning and knowledge representation. Other team performers will choose specific NI methods and develop algorithms for learning from them. These will be performers with a strong track record of learning related to the specific NI method they are targeting, yet able to propose innovative algorithmic ideas of how to efficiently perform learning given the inputs for their NI method. For learning system builders, the BL program offers the opportunity to perform groundbreaking work in a new research area. Because curricula will be built by separate performers, learning system builders will be able to concentrate a high fraction of their effort on research while still being able to perform rigorous experiments in several domains. Each NI method represents a separate research problem, with interesting possibilities for the interaction between methods as well.

Curriculum Builder Team

The curriculum builder team will provide several curricula per year for multiple domains. This will include the assumed prior knowledge for the domain as well as an interface for the student to a simulator for that domain.

DARPA will provide a set of areas of interest. Within those areas, domains can be proposed by potential performers, allowing them to take advantage of their existing resources. Because of the diversity of domains, a successful curriculum team will likely be composed of multiple organizations, each providing expertise required for the diverse curricula.

An important additional role of the curriculum development team will be the arbitration of issues related to the genetic knowledge (e.g., interaction language, NI methods, interlingua representation). These will be set definitively at the beginning of the program, based on the specifications that will accompany this document (see External Materials section below), and will only be changed after the start of the program as a result of demonstrated necessity. This is due to the fact that any changes will affect both learning teams and curriculum developers. Part of the purpose of this Request For Information is to solicit comments on these specifications, as will be detailed below.

Important Elements of the Program

As mentioned above, for this program to be successful, a number of different elements need to be in place before the development of program components can begin: a specification of the interaction language, a specification of the interlingua, a fixed set of natural instruction methods, and a set of domains in which ladders will be built. Each of these will be discussed in turn. Note that each of these elements corresponds to questions being asked in the "Request for Information (RFI)" section below.

Interaction Language

Bootstrapped learning with a human instructor is an AI-complete problem. It would require solutions for the various modalities used to interact with the student such as natural language understanding, computer vision, and diagram understanding. This is not the focus of the BL program. Instead, for the purposes of the program, a language of abstractions of the raw interactions that occur between humans during instruction will be termed the *interaction language*. The goal of this language is not to maximize expressivity. Rather, it is to provide a simplified representation capable of expressing materials communicated during a wide range of instructional methods.

In our ongoing investigations, the following interaction modalities have been identified:

- *Linguistic* – written and/or spoken signals to the student
- *World perception* – the student's perception of the state of the world, or of a hypothetical state of the world (e.g., that an instructor would show as an example)
- *World action* – base actions in the world. Actions that the student can take in the world when practicing, or that it can observe the instructor making in the world (e.g., in the NI method *By Demonstration*)
- *Gesture* – an instructor's pointing gestures to features in other modalities (such as pointing at an action or object in the world)
- *Diagrams* – visual abstractions used in instruction
- *Instructional cues* – Specific ways of conveying the structuring of training materials: the dependencies between lessons, the objective for a given lesson, etc.

The current specification of the interaction language will be given in an accompanying document (see External Materials section below). It is vital to the program that this language be both expressive enough to use for interesting domains, yet simple enough that it can be learned. A rule of thumb is a goal of 80% expressivity in this program, with the assumption that the relevant portion of the remaining 20% would be fleshed out in a specific practical application of bootstrapped learning.

Interlingua

As mentioned above, because lessons build on previous lessons, a common language for representing knowledge is necessary. As with the interaction language, to do this in general is probably a Knowledge Representation-complete problem. For the purposes of the BL program, something that is simple enough that it could be generated and refined by learning processes, yet still expressive enough to yield interesting problems is required. The following types of knowledge need to be represented:

- *Syntactic (ontological)* – domain objects and actions; function and predicate types including type restrictions on parameters and return values; etc.
- *Logical* – world knowledge and inference rules
- *Procedural* – knowledge of how to do things in the world
- *Functional* – knowledge of how to compute complex functions by composing smaller ones

The current specification of the interlingua will be given in an accompanying document (see External Materials section below). Note that, as for the interaction language, it is vital that the interlingua be expressive, yet simple enough to be learned. This also follows the 80% rule of thumb given for the interaction language.

Natural Instruction Methods

As discussed above, NI methods describe different ways of instruction. Some of the methods identified thus far include:

- *By Rote* – the instructor defines a concept linguistically, and the student delineates the space of where this potentially ambiguous concept could fit into its model of the world, and chooses an initial placement. This placement can be refined by later NI methods (e.g., *By Practice)*. Note that, for BL, *By Rote* does not connote repeated study for memorization, as it can in humans, but rather, the appropriate integration of that memorized knowledge.
- *By Annotated Example* – the instructor shows and comments on examples
- *By Demonstration* – the instructor demonstrates a procedure
- *By Practice* – the instructor gives the student a task and a function to maximize and lets the student practice to refine its behavior
- *By Refinement* – the instructor corrects student mistakes, helping the student to refine its knowledge

This is certainly not a complete list of NI methods. Prospective learning teams will propose their own set of NI methods to cover. All NI methods that will be used in the program must also have a specification in terms of which interaction modalities they use, and what interaction they expect.

Important characteristics of an NI method for the BL program include the following (note that any one of these characteristics is enough to make an NI method potentially interesting):

- *Efficiency* – a computationally efficient learning algorithm can be developed to learn from the interactions employed by this method
- *Naturalness* – an idealization of instruction methods that humans use
- *Usefulness* – could be used to instruct computing systems to perform a variety of tasks in a variety of domains
- *Practical* – this method is much simpler than any other method for instructing the same tasks (e.g. easier than simply programming the same task.).
- *Encapsulatability* – a central aim of the research agenda is to provide a testbed where new learning methods can be rapidly modified and retested. This ability is critical for the inherently empirical nature of this investigation. Thus the ideal NI methods must be restricted to where the instructional materials can be approximated by an entirely automated method. So if one were to propose using instructor feedback on student solutions, one would need to characterize the class of such interactions where some expert system is capable of delivering the needed feedback without benefit of human intervention.

Ladder Domains

Ideally, a domain/curriculum:

- Is inexpensive to build
    - Requires limited background knowledge
    - Leverages existing simulators and training materials
    - Is easy to provide the "same" content to humans (for the domain used for human comparison experiments)
- Provides complex instruction
    - As measured by an observed large increment in human performance after instruction
    - Multiple layers of (sub-)concepts and (sub-)procedures
    - Requires relational knowledge and representation shifts
- Resides in cyber domains where the perception problem is easier
- Represents "natural tasks" i.e., tasks currently taught or that could be taught to humans

An Example of Bootstrapped Learning

Consider learning in the domain of Unmanned Aerial Vehicle (UAV) surveillance. In a scenario where the human is doing the learning, the UAV has a human operator, whose job it is to fly the UAV (using remote control) to accomplish a set of predetermined tasks during the flight, given the relative importance of the tasks, restrictions in terms of flight distance, fuel, etc. Imagine deciding to train this human operator to not only try to accomplish the set tasks for the mission, but also to opportunistically look for other situations as well. This kind of training would probably involve (among other things), (1) the definitions of the various situations of interest, and (2) instruction on integrating these opportunistic goals with other priorities (e.g., when it is more important to take an opportunistic goal, even if it means no other goal will be accomplished in this mission).

Now consider the same situation, but this time with a bootstrapped learning system. Assume it is an autonomous agent with a planning and execution system, a computer vision system, etc., and, as was the case with the human controller, already has the ability to fly missions and accomplish various prioritized tasks within the UAV domain (injected knowledge).

Imagine that where the UAV is deployed, the monitoring of truck-to-truck (T2T) cargo transfers becomes of interest, as it is suspected that this is being used to smuggle goods into a secure location. The UAV controller (human or AI) must be taught to opportunistically survey truck-to-truck transfers it sees as it goes about its other tasks. This could be done in a curriculum ladder consisting of three rungs (lessons): recognition of T2T transfers; what to do when T2T transfers are observed; and refinement of priorities versus other goals. Each rung is described below.

Recognize T2T Transfers

The first rung is a lesson designed to teach the UAV controller to recognize when a T2T transfer is taking place. This might be taught using the NI method *By Annotated Example* (there are certainly other methods by which this might be taught). In this NI method, the instructor shows several snapshots of hypothetical states of the world (encoded in the *world perception* part of the interaction language for the AI UAV controller), comments on each example (encoded in the *linguistic* part of the interaction language), and points to salient features in the pictures (encoded in the *gesture* part of the interaction language). For example, one snapshot may be accompanied with the utterance "This is a truck-to-truck transfer since the rear of this truck [*points to a truck in image*] is close to the rear of this truck [*points to another other truck*]."

From this lesson, the learning process produces a new logical rule (encoded in the interlingua) that gives the conditions that would lead it to infer that a truck-to-truck transfer is occurring. In the case that the controller did not know beforehand the concept of "rear," the concept and function to compute it may have been learned in this lesson as well.

Opportunistic Rule

The next rung teaches the UAV controller to opportunistically survey a T2T transfer if it sees one. This might be taught using the NI method *By Rote*, where the instructor says, "If you see a truck-to-truck transfer, survey it." Note that the form of this utterance gives the rule to the student, more or less exactly as it is needed. The interesting task for the learner is knowing *where* to put this rule such that it will have the intended effect, as well as *which priority* to assign to the task of surveying the T2T transfer. Note that this rule *builds upon* the concept of T2T transfer which was learned in the previous rung.

Priority Refinement

The final rung is meant to teach the UAV controller to correctly prioritize tasks. This may be taught by the NI method *By Practice*, where the instructor has the UAV controller fly various missions and then scores performance based on whether or not recognized T2T transfers were actually surveyed. To show the importance of this, consider a scenario where the UAV controller initially (as the result of the previous rung) set a very low priority on surveying T2T transfers. In this case, it may report after a mission that it saw many T2T transfers, but decided not to survey them because it had more important tasks to do. The flip side of this is a controller that spent all day surveying T2T transfers, but didn't accomplish some very important task it was given for the mission.

Request For Information (RFI)

As discussed above, the interaction language, the NI method specification, and interlingua specification are vital requirements for the success of the BL program – both for curriculum developers and for learning teams. A seed effort is currently underway to define these specifications, and a draft will be available (see External Materials section below for details). This RFI has been released prior to those specifications in order to allow the community to consider the program as a whole, and to provide advance notice of the date of our workshop. DARPA/IPTO requests feedback from the scientific community on the following five aspects of the bootstrapped learning program proposal. All information contained in the RFI is preliminary and subject to modification; it is in no way binding on behalf of the Government.

1. Natural Instruction methods

What "Natural Instruction" methods should be the focus of this investigation? The goal is to identify a small set of very widely applicable and practical methods for instructing computing systems.

- What important NI methods are not listed? What makes these interesting for the BL program?
- What kinds of learning algorithms could be used for these (or other) NI methods?

2. Interaction language specification

What interaction modalities are most critical given the collection of NI methods DARPA wishes to develop? Furthermore, what is the *simplest* representation of the interaction that is still rich enough to support an interesting (and practically useful) set of learning tasks?

- Are the modalities listed sufficient for the goals of the program? Should others be added?
- For each modality, is there important information that is not represented? How can the specification be extended to represent it?
- Are the representations too difficult to perform learning with? Too easy?

3. The interlingua specification

The interlingua is the exchange format of knowledge learned between different learning methods. This is also the language of the “injected” knowledge (see above) used to specify the background knowledge that provides the starting point for the bootstrap learning process. Each learning process will have its own (possibly more expressive) internal representations. The interlingua needs to only be expressive enough to serve as an interchange language. Does it serve that purpose? In particular:

- Are these representations expressive enough? What important aspects are not shown and how could the interlingua be extended to represent them?
- Are they simple enough to *learn* and refine using instructor-guided learning processes?

4. Curricula suggestions

This RFI outlines characteristics of a good ladder domain and training curriculum. In the AAAI presentation (see External Materials section) there is a single slide outlining possible domain area choices. What specific task domains (simulators, test problems, training curricula) suggest themselves as ideal candidates for bootstrapped learning?

5. Intellectual property considerations

The curricula and learning systems developed as part of this program serve as an ideal testbed for ongoing research on instructable computing systems. Thus one possible future for the work developed in this program could be open forums or competitions (like RoboCup) where external investigators use curricula to develop their own learning algorithms. Progress might also be enhanced by allowing future work to build on learning components developed as part of this program, either by working with such methods (as a black box) or by extending or enhancing the methods themselves.

Obviously such work could only be done if DARPA had the rights (with some set of appropriate restrictions) to release materials developed during this program, and if those materials did not depend on components that were proprietary or required licenses.

One basis upon which potential performers will be judged is the IP license agreements acceptable to their organization. DARPA is asking the potential performer community about the impact of alternative IP license agreements and the likely restrictions that performers may require before allowing their materials to be consistent with this approach, both during the project and by users of project results. In particular:

- How might the software be creatively licensed, so that it could be broadly used for the purposes of BL-type research but still retain value to industry developer(s) for other purposes?
- Could the project be done entirely using open source applications and open architectures? What, if any, clear limitations exist in the research goals that might prevent a full open source implementation?
- Could a Creative Commons or similar licensing scheme, even if not optimal for private developers, still be sufficiently protective for this project? If not, why?

- Would developers of proprietary code allow source code to be provided for use in this kind of a learning-oriented project, as well as afterwards? For non-government users?
- Can realistic cost projections of licensing fees be projected for future users, if proprietary code is used? If so, for what periods of time into the future would such estimates remain valid? As much as 5 years after project completion?

Other comments/suggestions for the program

RFI Response and Participation Requirements

DARPA invites interested researchers and potential performers to respond in any or all of the following three ways: registering and providing written responses to this RFI, registering and participation in the RFI Workshop, and/or participation on a teaming page for the upcoming Bootstrapped Learning program.

Registering to Send a Written Response to the RFI and Workshop Statement of Interest:

Registration to respond to this RFI and submitting a Workshop Statement of Interest can be accomplished by completing the RFI Cover Sheet at: http://csc-ballston.dmeid.org/rfi/RFIindex.asp?RFIid=06-33. Once this Cover Sheet has been completed, an online confirmation sheet will appear. This confirmation sheet should be attached to the RFI Response, if any is submitted. See Section IV "Application and Submission Information" below for details on the specific content and form of the RFI Response.

(Note: The RFI Registration is separate from registration to attend the Workshop. See Workshop Registration details below.)

RFI Registration will include an optional Workshop Statement of Interest section. This section will ask for the following information:

- Brief background of the respondent.
- Brief description of interest in the program, including possible roles of respondent in the program (learning team integrator, learning process researcher, curriculum developer, other).
- Respondent's interest in attending the workshop.

While the RFI Registration does not entail a commitment to submit a written response to the RFI, or to attend the workshop, all interested parties should submit an RFI Registration no later than September 15, 2006. This registration and the optional Workshop Statement of Interest section will be used to assess participant interests/capabilities.

Participation in the RFI Workshop

A workshop will be held for interested parties on September 26, 2006, in Palo Alto, California. There will be an optional second day on September 27, 2006 for 10-15 minute sidebars with the program manager, and teaming discussions. An overview of the anticipated Bootstrapped Learning program will be presented, including both structured discussions of the RFI issues, unstructured discussion periods, and Q&A. If interested in attending, please visit the Registration Website at:

http://csc-ballston.dmeid.org/darpa/registration/intro.asp?regCode=cMA2NfJa

In order to attend the workshop, attendees ***must*** be registered by the September 20, 2006 deadline. DARPA reserves the right to limit attendance if participation becomes excessive. Because of space limitations, participation may be restricted to a single representative of a given institution and/or a single co-author of a given Statement of Interest. Non-U.S. citizens may attend pending the completion of the Foreign National Information Request Form (located on the registration site) and subsequent notification by DARPA of approval to attend. Attendance at the workshop is voluntary and is not required in order to submit to this RFI or any subsequent Broad Agency Announcements (BAA) or research solicitations on this topic. There is no fee for attending the workshop, and DARPA will not provide cost reimbursement for workshop attendance.

Participation on the Teaming Pages

Teaming Pages will be established for parties interested in teaming for an anticipated upcoming BAA. A separate Teaming Page will be provided for "Learning" (http://csc-ballston.dmeid.org/baa/BL-Learning_Teaming.htm) and for "Curriculum" (http://csc-ballston.dmeid.org/baa/BL-Curriculum_Teaming.htm). Note: all information submitted will be available to public viewing.

If you wish to be added to one of the teaming pages, please submit the following information to the BLSolicitations@darpa.mil:

- Name
- Organization
- E-mail Address
- Phone Number
- Fax Number
- Organization URL
- Denote Teaming Interest as one of the following Roles:
  (a) Team Lead (prime); (b) Team Contributor (sub); (c) or Either (willing to consider both)
- Identify if you wish to Team under (a) Learning or (b) Curriculum. Note: for purposes of the RFI, you may sign up for both team sites however, under the BAA, you will only be allowed to propose against one or the other.

- Identify Contributor Areas of Expertise.
  For Curricula developers: List specific task domains, simulators, etc.
  For Learning developers: List learning/acquisition areas of interest

OPTIONAL: Include a one-page teaming paper (in pdf format). This teaming paper is intended to highlight potential contributions to those forming teams for the program. See details below.

Learning Team Contributor
The optional one-page teaming paper should outline the type of contribution envisioned to a larger team:

- the type of natural instruction being proposed,
- the algorithmic approach is being considered,
- an argument for the practicality and generality of both the method and approach.

In addition, interested parties may submit a second optional section outlining relevant publications or other work to be highlighted.

Curriculum Team Contributor
Curriculum team contributors may also provide a one-page teaming paper. In their case, the teaming paper should outline the domain and software assets available.

External Materials

Other relevant information may be found at http://csc-ballston.dmeid.org/baa/BL_External_Materials.htm. Note: the information on this site will be updated on a regular basis so interested parties should check back often. The information posted will include:

- A Bootstrapped Learning briefing presented at the American Association for Artificial Intelligence (AAAI) held July 16–20, 2006.
- An initial draft of the interaction language discussed in this RFI will be available no later than September 11, 2006. ***Ideally, respondents to this RFI will wait to see the details of this draft interaction language prior to responding to the RFI.***

II. Award Information

This notice, which constitutes the complete RFI package, is not a Request for Proposals (RFP), and is not to be construed as a commitment by the Government to issue a solicitation or ultimately award a contract. Responses will not be considered as proposals nor will any award be made as a result of this solicitation. The Government is not interested in specific company capability information and will not entertain such submissions. Any costs incurred as a result of responding to this announcement shall be borne by the respondent and cannot be charged to the Government for reimbursement.

III. Eligibility Information

1. Eligible Applicants
All responsible sources capable of satisfying the Government's needs may submit a Statement of Interest or RFI Response that shall be considered by DARPA.

2. Cost Sharing or Matching – N/A

IV. Application and Submission Information

1. Address to Request Application Package
This announcement contains all information required to submit a statement of interest and RFI response. No additional forms, kits, or other materials are needed.

2. Content and Form of Application Submission

DARPA/IPTO requires completion of an online RFI Cover Sheet for each RFI response. This Cover Sheet is located at: http://csc-ballston.dmeid.org/rfi/rfiindex.asp?RFIid=06-33. ***Please note that if you have already registered for the RFI as noted above in the "RFI Response and Participation Requirements" section, you should not complete a second cover sheet. If you need to edit a previously entered cover sheet/registration form, you may do so by logging back in with your userid and password.***

RFI Responses are limited to 8 pages in length (not including cover sheet), and respondents are encouraged to use fewer pages, to be as succinct as possible while at the same time providing actionable insight. Each response should comprise the following sections: Section I. Cover Sheet: This should be the confirmation sheet referred to above under "Registering to Send a Written Response to the RFI and Workshop Statement of Interest". Section II. Expand on the background of the respondent as listed on the Cover Sheet. Section III. Address any or all of the sections listed above in the "Request for Information" portion of this document. Each insight regarding the posed questions should be stated as a *one or two sentence summary* followed by a paragraph elaboration or explanation of the stated insight. (Multiple insights per question are permissible.) Section IV. Expanded description of interest in the program, including possible roles of respondent in the program (learning team integrator, learning process researcher, curriculum developer, other). Section V. (Optional) Description of a natural instruction method for which the responder would be interested in building a learning process. Section VI. Any additional information.

It is anticipated that participants will not provide proprietary information given the open nature of the workshop. However, if any propriety material is submitted, each section must be clearly marked as such. Any information not clearly marked as proprietary will be considered to be public information. All submissions of any proprietary information may be handled by non-government personnel bound by nondisclosure agreements. This RFI incorporates by reference FAR 52.215-3, "Request for Information or Solicitation for Planning Purposes (OCT 1997)," with the same force and effect as if it were given in full text [reference paragraph (c) of this provision, the "purpose" of this RFI is detailed in this announcement].

Respondents must submit two paper copies of the full response and one electronic copy (in Microsoft Word or Adobe PDF on a CD ROM). Disks must be clearly labeled with RFI 06-33, respondent organization, and points of contact. The RFI responses must be submitted to: DARPA/IPTO, Attn: RFI 06-33, 3701 N. Fairfax Drive, Arlington, VA 22203-1714.

3. Submission Dates and Times
RFI Registration and optional Statements of Interest must be received online at http://csc-ballston.dmeid.org/rfi/rfiindex.asp?RFIid=06-33 by 12:00 NOON (ET), September 15, 2006.

RFI responses will be considered if they are received at DARPA by 12:00 NOON (ET), September 20, 2006. DARPA will acknowledge receipt of RFI responses via email and assign control numbers that should be used in all further correspondence regarding the submissions.

4. Intergovernmental Review – N/A

5. Funding Restrictions - N/A

6. Other Submission Requirements – N/A

V. Application Review Information

1. Criteria – N/A

2. Review Process
Although DARPA will acknowledge receipt of RFI responses, no feedback will be provided.

VI. Award Information Administration – N/A

1. Award Notices – N/A

2. Administrative and National Policy Requirements – N/A

3. Reporting – N/A

VII. Agency Contacts
All administrative correspondence and questions concerning this announcement should be directed to one of the following administrative addresses:

Fax: 703-741-7804, Addressed to: DARPA/IPTO, Attn: RFI 06-33
Electronic Mail: BLSolicitations@darpa.mil
Electronic File Retrieval: http://www.darpa.mil/ipto/Solicitations/solicitations.htm

Mail to: DARPA/IPTO
ATTN: RFI 06-33
3701 N. Fairfax Drive
Arlington, VA 22203-1714